



036

#4

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : Loren Christensen

Serial No.: 09/843,471

Filed: April 26, 2001

Title: HIGH PERFORMANCE DISTRIBUTED DISCOVERY SYSTEM

Docket No.: 33556

LETTER

Commissioner for Patents
Washington, D.C. 20231

Sir:

Enclosed is a certified copy of Canadian Patent Application No. 2,322,117; the priority of which has been claimed in the above-identified application.

Respectfully submitted,

PEARNE & GORDON LLP

By: John P. Murtaugh
John P. Murtaugh, Reg. No. 34226

526 Superior Avenue East
Suite 1200
Cleveland, Ohio 44114-1484
(216) 579-1700

Date: 7-16-01

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, Washington, D.C. 20231 on the date indicated below.

John P. Murtaugh

Name of Attorney for Applicant(s)

July 16, 2001 John P. Murtaugh
Date Signature of Attorney



Offic de la propriété
int lctuelle
du Canada

Canadian
Int lctual Property
Office

Un organisme
d'Industrie Canada

An Agency of
Industry Canada



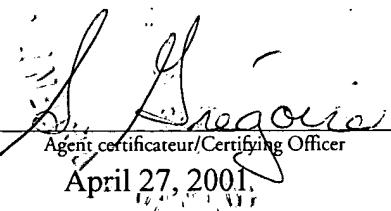
Bureau canadien
des brevets
Certification

La présente atteste que les documents
ci-joints, dont la liste figure ci-dessous,
sont des copies authentiques des docu-
ments déposés au Bureau des brevets.

Canadian Patent
Office
Certification

This is to certify that the documents
attached hereto and identified below are
true copies of the documents on file in
the Patent Office.

Specification and Drawing, as originally filed, with Application for Patent Serial No:
2,322,117, on October 3, 2000, by LINMOR TECHNOLOGIES INC., assignee of
Loren Christensen, for "High Performance Distributed Discovery System".


Agent certificateur/Certifying Officer

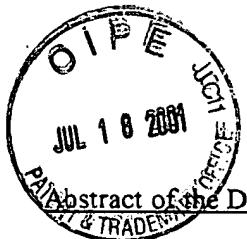
April 27, 2001

Date

Canada

(CIPO 68)
01-12-00

O P I C  C I P O



A high performance distributed discovery system for the discovery of the network topology
of devices of a high speed data network comprising a plurality of discovery engines on at least
5 one and preferably a plurality of data collection node computers for the polling and
registration of discovered objects into a distributed topology database whereby parallel
processing is provided. The topology database is periodically pulled down to at least one
performance monitoring server computer to provide for network management.

10

15

20

High Performance Distributed Discovery System

Inventor: Loren Christensen

5

Field of the Invention

The present invention relates to the discovery of the network topology of devices comprising a high speed data network, and more particularly to a high performance distributed discovery system.

Background of the Invention

Today's high speed data networks contain an ever-growing number of devices. A network needs to be monitored for the existence, disappearance, reappearance and status of traditional network devices such as routers, hubs and bridges and more recently high speed switching devices such as ATM, Frame Relay, DSL, VoIP and Cable Modems.

In order to enable network monitoring, a process known as discovery is typically performed. Discovery is the process by which network management systems selectively poll a network to discover very large numbers of objects in a very short period of time, without introducing excessive network traffic. It is the function of a discovery system to discover devices on a network and the structure of that network. Discovery is primarily intended to get network management users quickly up to speed, track changes in the network, update network maps and report on these changes.

Discovery typically further involves discovering the configuration of individual devices, their relationship, as well as discovering interconnection links or implied relationships.

30

5 In the past, rapid discovery was not an issue, since the level of scalability of performance monitoring did not require the depth of discovery that is now required. Major advances in scalability have recently been achieved in performance monitoring and as performance monitoring scales to manage larger and larger networks the scalability of discovery must advance accordingly in order to deal with the inevitable increase in the number of network objects and react quickly to changes in network topology.

10 At present network devices are typically polled over long distances from the network management system. This consumes valuable bandwidth and results in increased processing times and potential data loss. As well, customers often dislike inadvertent access around their firewalls, via the common connection to the network performance monitoring server computer. Therefore, what is needed is a method of object discovery that is proximal to the managed network.

15 For the foregoing reasons, there is a need for an economical method of network topology discovery that provides for high speed polling, high object capacity, scalability, and proximity to managed networks, while preserving security policies that are inherent in the network domain configuration.

20 Summary of the Invention

25 The present invention is directed to a high performance distributed discovery system that satisfies this need. The system, leveraging the functionality of a high speed communications network, comprises distributing records of discovered network devices using a plurality of discovery engine instances located on at least one data collection node computer whereby the resulting distributed record compilation comprises a distributed network topology database.

The distributed network topology database is accessed using at least one performance monitor server computer to facilitate network management. At least one discovery engine instance is located on the data collection node computers on a ratio of one engine instance to one central processing unit whereby the total number of engine instances is at least two so as to enable the parallel processing of the distributed network topology database.

In aspects of the invention a vendor specific discovery subroutine is launched upon detection by the system of a non-MIB II standard device so as to query the vendor's private MIB using a vendor specific algorithm.

Advances in overall scalability are achieved by dividing the workload of network topology discovery across several computing nodes. The discovery job is distributed across all the data collectors such that the only requirement for each data collector is to be able to reach, typically via TCP/IP and SNMP, the nodes and networks for which it is responsible. This reachability requirement already exists for telemetry, in any case, and has therefore already been provided for.

Brief Description of the Drawings

These and other features, aspects, and advantages of the present invention will become better understood with regard to the following description, appended claims, and accompanying drawings where:

Figure 1 is an overview of the high performance distributed discovery system.

Detailed Description of the Presently Preferred Embodiment

As shown in figure 1, the high performance distributed discovery system comprises at least one data collection (DC) node computer 12 for polling and registering managed network 14 objects whereby the resulting distributed record compilation comprises a

distributed topology database 16 and at least one performance monitor (PM) server computer 18 in network contact with the DC node computers 12 for accessing the distributed topology database 16.

5 A plurality of discovery engine instances 20 are located on the DC node computers 12 on a ratio of one engine instance 20 to one central processing unit so as to provide for the parallel processing of the distributed topology database 16.

10 The discovery engine 20 is comprised of a base program and a scalable family of vendor-specific discovery subroutines. The base program is designed to query and register any IP device and subsequently obtain detailed device, state and topology information for any IP device that responds to an SNMP query, such as any device that is managed by an SNMP agent. The base program discovers detailed information for any device that supports the standard MIB-II, but not the vendor's private MIB. The discovery of detailed information 15 from a vendor's private MIB is accomplished through what is known as vendor-specific discovery subroutines.

20 The DC node computers are responsible for telemetry to the managed elements and management of the topology database. The PM server computers provide system control and reporting interface.

The vendor specific discovery subroutines extend the base discovery application to provide for inter-operability with a multiplicity of ATM and FR vendors' equipment.

25 These discovery subroutines are lightweight independent applications that are launched whenever the main discovery program detects a particular vendor's hardware. The discovery subroutines contain vendor-specific algorithms designed to query the vendor's private management information base (MIB).

Launch points for each discovery subroutine are included in the main program. So, if during the normal operation of discovery a valid element value is encountered identifying a specific vendor's hardware, the appropriate discovery subroutine is launched.

5 All of the processing intensive data collection takes place as close to the customers network and network devices as possible, thereby providing for faster discovery as well as distributed storage and processing.

10 The proximal topology of the DC node computers in relation to the managed network provides for inherent scalability and a reduction in required bandwidth. As well, the ability to utilize excess memory and disk storage resources on the DC node computers facilitates the discovery of larger networks. The aggregate resources of many DC node computers is far greater than that available on any one PM server computer. Advances in overall scalability are achieved by dividing the workload of network topology discovery across several computing
15 nodes. The discovery job is distributed across all the data collectors such that the only requirement for each data collector is to be able to reach, typically via TCP/IP and SNMP, the nodes and networks for which it is responsible. This reachability requirement already exists for telemetry, in any case, and has therefore already been provided for.

20 All the discovery and topology database storage is taking place behind the client's firewall requiring only a minimal amount of management traffic to be exerted on the network to generate reports. PM server computers are utilized to access the distributed topology database for object management.

25 The performance monitor was unwittingly becoming a router. This invention removes this unwanted side-effect thereby enhancing security.

Devices are reliably re-discovered, thereby enabling the tracking of changes to a network's topology as it evolves in real time or near real time.

The ability to limit what is discovered by criteria such as vendor & device type has been added thereby eliminating the need to specify the address of each device when discovering the network.

5 The system will not re-discover existing devices unless explicitly requested to do so, which is significant when discovering a large network that is typically discovered in stages.

 The system handles timeouts in a more reliable manner. This is important on wide area networks where timeouts are more common during discovery.

10

 Since all the discovery sub-tasks can be performed simultaneously, the overall time to characterize the customer's network is reduced. This enables discovery to deal with larger networks in a faster manner, and eliminates the performance monitor's reachability requirement with respect to managed elements.

15

 This will allow Network Service Providers to automatically discover more of the existing devices in their networks, permitting customers to reconcile what is really out in their network with what their administrative records tell them is out there. It has been shown that such verification can potentially lead to great cost savings in operations, as well as vastly improved discovery times as speed will now be directly correlated with the number of DC node computers deployed.

20

 In aspects of the invention unique algorithms selectively discover network devices based on "clues" picked up from existing information such as router tables and customer input.

25

 All resulting in the rapid automatic mapping of a customer's network for the purpose of object management down to unprecedently fine levels of granularity.

What is claimed is:

1. A network topology distributed discovery system, leveraging the functionality of a high speed communications network, comprising the steps of:

5 (i) distributing records of discovered network devices using a plurality of discovery engine instances located on at least one data collection node computer whereby the resulting distributed record compilation comprises a distributed network topology database; and
 (ii) importing the distributed network topology database onto at least one
10 performance monitor server computer so as to enable network management.

2. The system according to claim 1, wherein at least one discovery engine instance is located on the data collection node computers on a ratio of one engine instance to one central processing unit whereby the total number of engine instances is at least two so as to enable
15 the parallel processing of the distributed network topology database.

3. The system according to claim 1, wherein a vendor specific discovery subroutine is launched upon detection by the system of a non-MIB II standard device so as to query the vendor's private MIB using a vendor specific algorithm.

20 4. The system according to claim 1, wherein at least one performance monitor client computer is connected to the network so as to communicate remotely with the performance monitor server computers.

25 5. A network topology distributed discovery system, leveraging the functionality of a high speed communications network, comprising:

 (i) at least one data collection node computer connected to the network for discovering network devices using a plurality of discovery engine instances whereby a distributed network topology database is created; and

(ii) at least one performance monitor server computer having imported the distributed network topology database whereby network management is enabled.

6. The system according to claim 5, wherein at least one discovery engine instance is located
5 on the data collection node computers on a ratio of one engine instance to one central processing unit whereby the total number of engine instances for the system is at least two so as to enable the parallel processing of the network topology database.

7. The system according to claim 5, wherein a vendor specific discovery subroutine is
10 launched upon detection by the system of a non-MIB II standard device so as to query the vendor's private MIB using a vendor specific algorithm.

8. The system according to claim 5, wherein at least one performance monitor client computer
15 is connected to the network so as to communicate remotely with the performance monitor server computers.

9. A storage medium readable by an install server computer in a network topology distributed discovery system including the install server, leveraging the functionality of a high speed communications network, the storage medium encoding a computer process comprising:
20 (i) a processing portion for distributing records of discovered network devices using a plurality of discovery engine instances located on at least one data collection node computer whereby the resulting distributed record compilation comprises a distributed network topology database; and
(ii) a processing portion for importing the distributed network topology database onto
25 at least one performance monitor server computer so as to enable network management.

10. The system according to claim 9, wherein at least one discovery engine instance is located on the data collection node computers on a ratio of one engine instance to one central

processing unit whereby the total number of engine instances is at least two so as to enable the parallel processing of the network topology database.

11. The system according to claim 9, wherein a vendor specific discovery subroutine is
5 launched upon detection by the system of a non-MIB II standard device so as to query the vendor's private MIB using a vendor specific algorithm.

12. The system according to claim 9, wherein at least one performance monitor client computer is connected to the network so as to communicate remotely with the performance
10 monitor server computers.

1/1

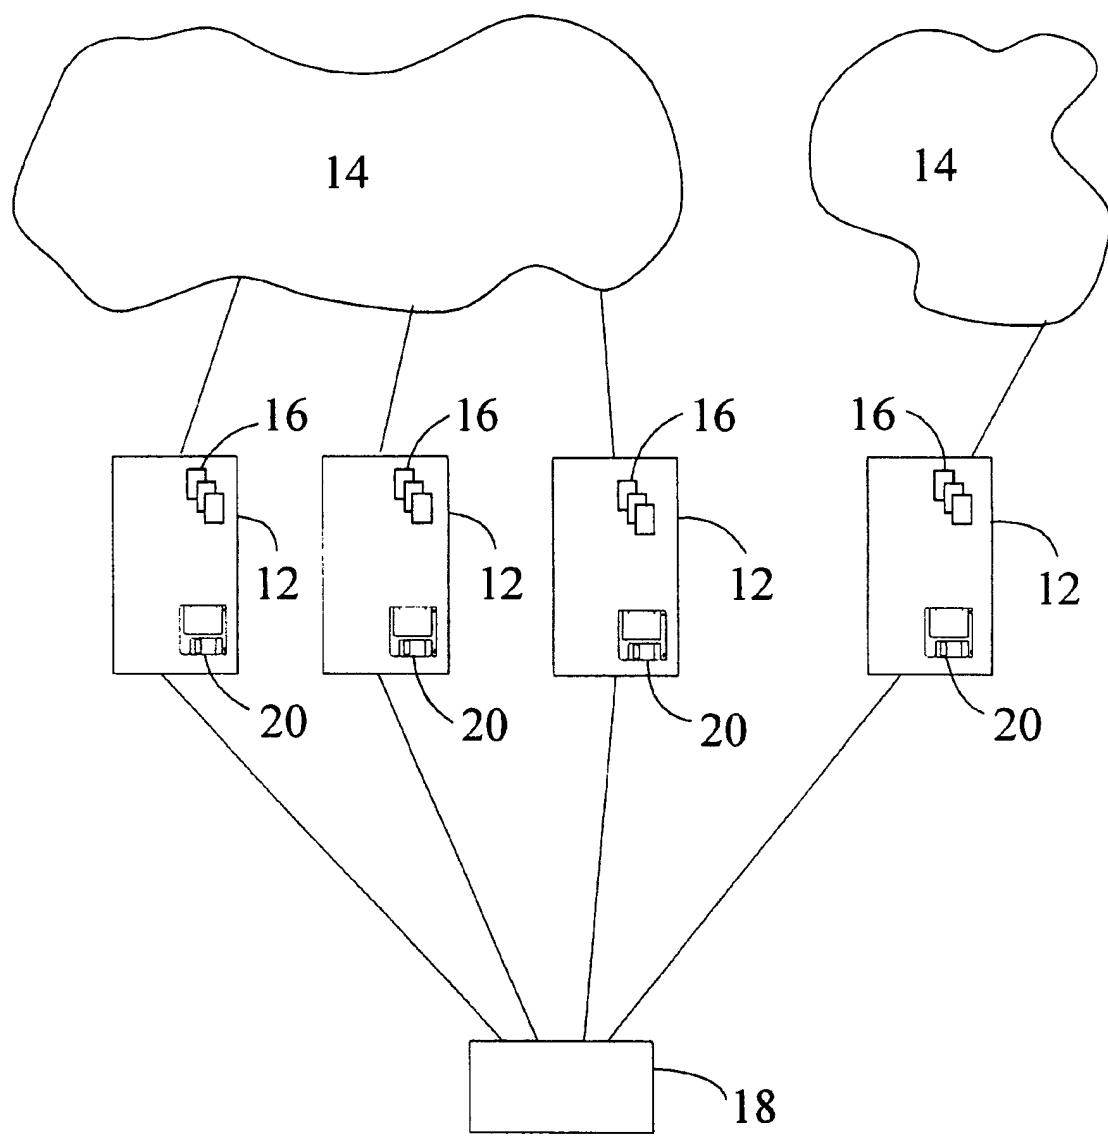


Figure 1

